

# Vision-Guided Flight Stability and Autonomy for Micro Air Vehicles

Scott M. Ettinger<sup>1,2</sup> Michael C. Nechyba<sup>1</sup> Peter G. Ifju<sup>2</sup> Martin Waszak<sup>3</sup>  
 sme@aemes.aero.ufl.edu nechyba@mil.ufl.edu pgi@aero.ufl.edu m.r.waszak@larc.nasa.gov

<sup>1</sup>Department of Electrical and Computer Engineering  
 University of Florida, Gainesville, FL 32611-6200

<sup>2</sup>Department of Aerospace Engineering, Mechanics and Engineering  
 University of Florida, Gainesville, FL 32611-6250

<sup>3</sup>Dynamics and Control Branch, NASA Langley Research Center  
 MS 132, Hampton, VA 23681-2199

## Abstract

*Substantial progress has been made recently towards designing, building and test-flying remotely piloted Micro Air Vehicles (MAVs). We seek to complement this progress in overcoming the aerodynamic obstacles to flight at very small scales with a vision-guided flight stability and autonomy system. The developed system is based on a robust horizon detection algorithm which we discuss in greater detail in a companion paper [5]. In this paper, we first motivate the use of computer vision for MAV autonomy, arguing that given current sensor technology, vision may be the only practical approach to the problem. We then briefly review our statistical vision-based horizon detection algorithm, which has been demonstrated at 30Hz with over 99.9% correct horizon identification. Next, we develop robust schemes for the detection of extreme MAV attitudes, where no horizon is visible, and for the detection of horizon estimation errors, due to external factors such as video transmission noise. Finally, we discuss our feedback controller for self-stabilized flight, and report results on vision-based autonomous flights of duration exceeding ten minutes.*

## 1. Introduction

Ever since humankind's first powered flight, research efforts have continually pushed the envelope to create flying machines that are faster and/or larger than ever before. Now, however, there is an effort to design aircraft at the other, largely unexplored end of the spectrum, where the desire for portable, low-altitude aerial surveillance has driven the development and testing of aircraft that are as small and slow as the laws of aerodynamics will permit — in other words, on the scale and in the operational range of small birds. Vehicles in this class of small-scale aircraft are known as *Micro Air Vehicles* or *MAVs*.

Equipped with small video cameras and transmitters, MAVs have great potential for surveillance and monitoring tasks in areas either too remote or too dangerous to send human scouts. Operational MAVs will enable a number of important missions, including chemical/radiation spill monitoring, forest-fire reconnaissance, visual monitoring of volcanic activity, surveys of natural disaster areas, and even inexpensive traffic and accident monitoring. Additional on-board sensors can further augment MAV mission profiles to include for example airborne chemical analysis.

In the military, one of the primary roles for MAVs will be as small-unit battlefield surveillance agents, where MAVs can act as an extended set of eyes in the sky for military units in the field. This use of MAV technology is intended to reduce the risk to military personnel and has, perhaps, taken on increased importance in light of the U.S.'s new war on terrorism, where special operations forces will play a crucial role. Virtually undetectable from the ground, MAVs could penetrate potential terrorist camps and other targets prior to any action against those targets, significantly raising the chance for overall mission success.

Researchers in the Aerospace Engineering Department at the University of Florida have established a long track record in designing, building and test-flying (remotely human-piloted) practical MAVs [7-9,14,15]. While much progress has been made in the design of ever smaller MAVs by researchers at UF and others in the past five years, significantly less progress has been made towards equipping these MAVs with autonomous capabilities that could significantly enhance the utility of MAVs for a wide array of missions.

The first step in achieving such MAV autonomy is basic stability and control. Here, we present such a *flight stability and control system*, based on vision processing of video from a camera on-board our MAVs. In this paper, we first motivate the use of computer vision for such a control system, and briefly review our related work in *vision-based horizon detection* which is discussed more fully in a companion paper [5], and forms the basis of the flight stability system presented here. Next, we address real-time control issues in the flight stability system, including extreme attitude detection (i.e. no horizon in the image), confidence measures for the detected horizon estimates, and filtering of horizon estimates over time. Finally we report some results of self-stabilized MAV flights over the campus of the University of Florida and over Fort Campbell, Kentucky.

## 2. Horizon detection

Below, we motivate vision-based approaches to the flight stability and control problem for MAVs and review related work in in vision-based horizon detection.

### 2.1 Motivation

MAV flight stability and control presents some difficult challenges. The low moments of inertia of MAVs make them vulnerable to rapid angular accelerations, a problem further complicated

by the fact that aerodynamic damping of angular rates decreases with a reduction in wingspan. Another potential source of instability for MAVs is the relative magnitudes of wind gusts, which are much higher at the MAV scale than for larger aircraft. In fact, wind gusts can typically be equal to or greater than the forward airspeed of the MAV itself. Thus, an average wind gust can immediately affect a dramatic change in the flight path of these vehicles.

Birds, the biological counterpart of mechanical MAVs, can offer some important insights into how one may best be able to overcome these problems. For aerospace and robotic researchers, the extraordinary capabilities of birds are a source of wonderment and frustration at the same time — wonderment, because birds exhibit a remarkably complex and rich set of behaviors, frustration, because the duplication of those behaviors in man-made systems has thus far been elusive. Given this sad state of affairs, it is natural for engineers and researchers to want to learn from and emulate these biological systems. Here, we do not intend to seek *structural* equivalence between artificial MAV systems and biological neural systems; rather, we seek to learn important *functional* lessons from biology. In studying the nervous system of birds, one basic observation holds true for virtually all of the thousands of different bird species: *Birds rely heavily on sharp eyes and vision to guide almost every aspect of their behavior* [5,16].

Biological systems, while forceful evidence of the importance of vision in flight, do not, however, in and of themselves warrant a computer-vision based approach to MAV autonomy. Other equally important factors guide this decision as well. Perhaps most critical, the technologies used in rate and acceleration sensors on larger aircraft are not currently available at the MAV scale. It has proven very difficult, if not impossible, to scale these technologies down to meet the very low payload requirements of MAVs. While a number of sensor technologies do currently exist in small enough packages to be used in MAV systems, these small sensors have sacrificed accuracy for reduced size and weight. Take, for example, MEMs (Micro Electro-Mechanical Systems) rate gyros and accelerometers. MEMs piezoelectric gyros, while only weighing approximately one gram, have drift rates on the order of  $100^\circ$  per minute and are highly sensitive to changes in temperature. While elaborate temperature calibration procedures can improve their accuracy somewhat, their use in inertial navigation is problematic at best.

Even if sufficient rate and acceleration sensors did exist, however, their use on MAVs may still not be the best allocation of payload capacity. For many potential MAV missions, vision may be the only practical sensor that can achieve required and/or desirable autonomous behaviors. Furthermore, given that surveillance has been identified as one their primary missions, MAVs must necessarily be equipped with on-board imaging sensors, such as cameras or infrared arrays. Thus, computer-vision techniques exploit already present sensors, rich in information content, to significantly extend the capabilities of MAVs, without increasing the MAV's required payload.

## 2.2 Vision-based horizon detection

Fundamentally, flight stability and control requires measurement of the MAV's angular orientation. While for larger aircraft this is typically estimated through the integration of the aircraft's

angular rates or accelerations, a vision-based system can directly measure the aircraft's orientation with respect to the ground. The two degrees of freedom critical for stability — the *bank angle*  $\phi$  and the *pitch angle*  $\theta$ <sup>1</sup> — can be derived from a line corresponding to the horizon as seen from a forward facing camera on the aircraft. In [5], we develop and report results for a vision-based horizon-detection algorithm, which attempts to minimize the intra-class variance of the ground and sky pixel distributions. Figure 1 illustrates several examples of our current algorithm at work. In each image, the yellow line indicates the algorithm's estimated location of the horizon. The algorithm has been demonstrated to run at 30 Hz on a 900 MHz x86 processor with image resolution of  $320 \times 240$ , although acceptable performance has been shown for image resolutions as low  $160 \times 120$ .

At different times of the day, and under both fair and cloudy conditions, we have gathered hours of video on-board our MAV, flying under manual control over terrain that includes roads, buildings large and small, meadows, wooded areas, and a lake. For these data, our horizon-detection algorithm correctly identifies the horizon in over 99.9% of cases.

## 3. Flight stability and control

In this section, we extend the basic horizon-detection algorithm to *real-time horizon tracking*. Below, we consider the following important issues: (1) extreme attitude detection, (2) error detection in horizon estimation, (3) filtering of the horizon estimate over time, and (4) basic feedback control and stabilization of the MAV.

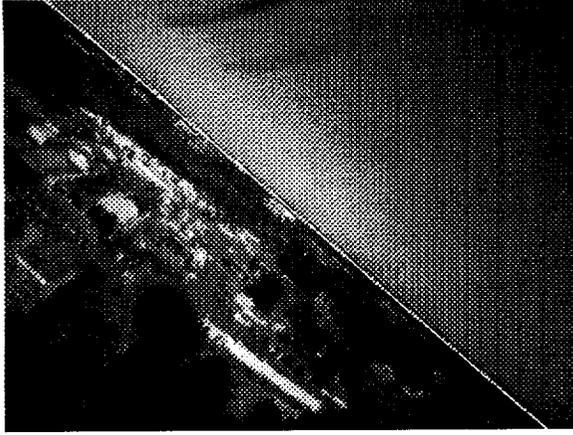
### 3.1 Extreme attitude detection

One of the implicit assumptions of the horizon detection algorithm is that there will always be a horizon in the images from the forward looking camera on board the MAV. In real-time control of the MAV, the MAV may, however, encounter times when no visible horizon appears in the image, if, for example, a gust of wind forces the nose of the aircraft too far up or down. Such cases cannot simply be ignored; if the aircraft is heading straight towards the ground, no horizon will be visible in the camera image, yet the control system will certainly be required to take action to save the MAV from certain and possibly catastrophic crashing.

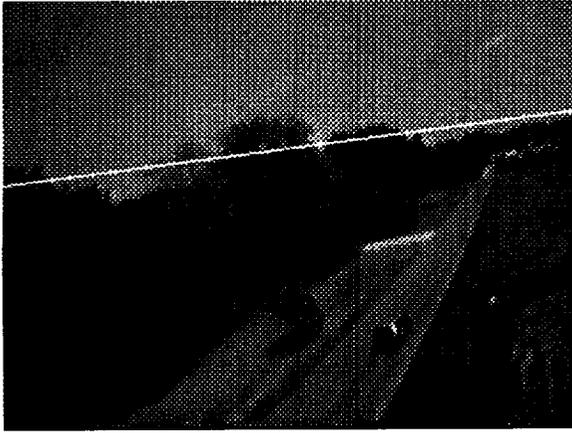
It is desired then, to be able to detect instances when the horizon is not in view of the camera, and if so to determine what action to take in order to bring the horizon back into view. There are two valuable sources of information which we can draw on to detect these types of extreme attitudes: (1) recent appearance of the sky and ground from previous time steps, and (2) recent location of the horizon line from previous time steps. For example, if the horizon line was recently estimated to lie near the top of the image, it is logical that a subsequent image without a horizon line is most likely a view of the ground. We can use these two pieces of information to quantitatively determine if the horizon line exists in the image and if not, to determine whether we are looking at the sky or the ground.

---

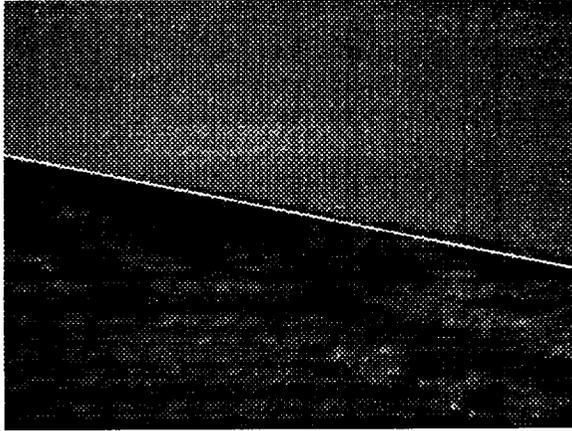
1. Instead of the pitch angle  $\theta$ , we actually recover the closely related pitch percentage  $\sigma$  in our horizon-detection algorithm.



(a)



(b)



(c)

**Fig. 1:** (a) Horizon detection example; (b) horizon detection with uneven horizon; (c) horizon detection with severe video interference.

As part of our horizon detection algorithm, we already compute the aggregate statistical properties of the sky and ground in RGB color space. Let us denote all the sky pixels in an image as,

$$x_i^s = [r_i^s \ g_i^s \ b_i^s], \quad i \in \{1, \dots, n_s\}, \quad (1)$$

where  $r_i^s$  denotes the red channel value,  $g_i^s$  denotes the green channel value and  $b_i^s$  denotes the blue channel value of the  $i$ th sky pixel, and let us denote all the ground pixels in an image as,

$$x_i^g = [r_i^g \ g_i^g \ b_i^g], \quad i \in \{1, \dots, n_g\}, \quad (2)$$

where  $r_i^g$  denotes the red channel value,  $g_i^g$  denotes the green channel value and  $b_i^g$  denotes the blue channel value of the  $i$ th ground pixel. Then, as part of the horizon detection algorithm, we compute the following statistical descriptions of the sky and ground appearance:

$$\mu_s = \frac{1}{n_s} \sum_{i=1}^{n_s} x_i^s, \quad \mu_g = \frac{1}{n_g} \sum_{i=1}^{n_g} x_i^g \quad (3)$$

$$\Sigma_s = \frac{1}{(n_s-1)} \sum_{i=1}^{n_s} (x_i^s - \mu_s)(x_i^s - \mu_s)^T \quad (4)$$

$$\Sigma_g = \frac{1}{(n_g-1)} \sum_{i=1}^{n_g} (x_i^g - \mu_g)(x_i^g - \mu_g)^T \quad (5)$$

where,  $\mu_s$  and  $\mu_g$  denote the means of the sky and ground pixels, respectively, and  $\Sigma_s$  and  $\Sigma_g$  denote the covariances of the sky and ground pixels, respectively.

We now use the above statistics to model the appearance of the sky and ground over a recent time history of the MAV's flight. Our general approach for detection of extreme attitudes keeps running statistical models for both the sky and ground from previous frames, where horizon lines were detected with a high degree of confidence. With each new frame, the result of the horizon detection algorithm can be checked by comparing the sky and ground models for the current frame with the computed, time-dependent statistical models for sky and ground. If the distributions on either side of the line in the current frame both appear to be more similar to the known ground distribution, then it would appear that the aircraft is pointing towards the ground. Conversely, if they both match the sky better, then it is advisable to nose downward. Interestingly, if the sky in the current frame matches the ground model while the ground in the current frame matches better with the sky model, we can detect situations where the plane is flying upside down.

One additional piece of information is required to implement the extreme attitude detection scheme, namely, a time history of the horizon line estimate. For the purposes of detecting extreme attitudes, we are most concerned with a recent history of the pitch percentage  $\sigma$ , the percentage of the image below the horizon line. One measure of that history is a running average  $\sigma_{avg}$  of the pitch percentage over the previous ten frames.

Upon startup of the system, the camera is assumed to be oriented such that the horizon is in its view. When the first frame of video is processed by the system, the means and covariance matrices of the ground and sky models are set equal to those found

by the horizon detection algorithm. The system then begins to update the models using the results of the horizon detection algorithm for a set number of initialization frames. Our current implementation uses 100 initialization frames (3.3 seconds). Once boot-strapped, it is necessary to continually update the sky and ground models as the aircraft flies to account for changes in lighting associated with changes in orientation and changes in landscape, etc. The running statistical models are updated as follows:

$$\Sigma_{s(t)} = \alpha \Sigma_{s(t)} + (1 - \alpha) \Sigma_s \quad (6)$$

$$\Sigma_{g(t)} = \alpha \Sigma_{g(t)} + (1 - \alpha) \Sigma_g \quad (7)$$

$$\mu_{s(t)} = \alpha \mu_{s(t)} + (1 - \alpha) \mu_s \quad (8)$$

$$\mu_{g(t)} = \alpha \mu_{g(t)} + (1 - \alpha) \mu_g \quad (9)$$

where  $\Sigma_{s(t)}$ ,  $\Sigma_{g(t)}$ ,  $\mu_{s(t)}$  and  $\mu_{g(t)}$  are the time-dependent model covariances and means, respectively, while  $\Sigma_s$ ,  $\Sigma_g$ ,  $\mu_s$  and  $\mu_g$  are the covariances and means for the current frame. Note that the constant  $\alpha$  controls how rapidly the models change over time.

For a new image, we first compute the estimated horizon for that image. We then compare the resultant current statistics with the running statistical models from previous frames, using the following four distance measures:

$$D_1 = (\mu_s - \mu_{s(t)})^T \Sigma_{s(t)}^{-1} (\mu_s - \mu_{s(t)}) + (\mu_s - \mu_{s(t)})^T \Sigma_s^{-1} (\mu_s - \mu_{s(t)}) \quad (10)$$

$$D_2 = (\mu_s - \mu_{g(t)})^T \Sigma_{g(t)}^{-1} (\mu_s - \mu_{g(t)}) + (\mu_s - \mu_{g(t)})^T \Sigma_g^{-1} (\mu_s - \mu_{g(t)}) \quad (11)$$

$$D_3 = (\mu_g - \mu_{s(t)})^T \Sigma_{s(t)}^{-1} (\mu_g - \mu_{s(t)}) + (\mu_g - \mu_{s(t)})^T \Sigma_s^{-1} (\mu_g - \mu_{s(t)}) \quad (12)$$

$$D_4 = (\mu_g - \mu_{g(t)})^T \Sigma_{g(t)}^{-1} (\mu_g - \mu_{g(t)}) + (\mu_g - \mu_{g(t)})^T \Sigma_g^{-1} (\mu_g - \mu_{g(t)}) \quad (13)$$

The value of  $D_1$  measures the similarity between the region selected as the sky by the horizon detection algorithm in the current frame and the sky model from recent frames.  $D_2$  represents the similarity between the currently computed sky region and the ground model from recent frames. Likewise, the values of  $D_3$  and  $D_4$  are the similarity measures between the current ground region and the sky and ground models from recent frames, respectively. Table 1 now summarizes four possible cases and the conclusions we can draw for each case.

The determinations in the above table can now be combined with the past history of the horizon line to decide what action to take. If the current frame is determined to be normal by the validity test (case 1), then the horizon estimate is assumed to be accurate, and commands sent to the MAV are determined by the normal control system loop described in Section 3.4. Also, the statistics of the validated frame are used to update the sky and ground models per equations (6) through (9). If the validity test returns a higher likelihood of all ground (case 2), we verify that result with the recent history of the horizon line  $\sigma_{avg}$  to deter-

**Table 1: Extreme attitude detection**

<i>case</i>	<i>condition</i>	<i>conclusion</i>
1	$D_1 < D_2$ and $D_3 > D_4$	valid horizon present
2	$D_1 > D_2$ and $D_3 > D_4$	all ground
3	$D_1 < D_2$ and $D_3 < D_4$	all sky
4	$D_1 > D_2$ and $D_3 < D_4$	upside down

mine what action to take. When the value of  $\sigma_{avg}$  is above a set threshold, then the system goes into a “pull-up” mode that sends commands to the aircraft to rapidly increase its pitch angle. A value of 0.8 was used for this threshold. While the system is in pull-up mode, the time-dependent statistical models are not updated since the horizon estimate during this time will most likely be incorrect. Also during pull-up mode,  $\sigma_{avg}$  is only updated with the estimated value of  $\sigma$  if the validity test indicates the current frame has returned to a visible horizon line; otherwise,  $\sigma_{avg}$  is updated using a value of 1.01. The system will stay in pull up mode until a valid horizon is detected. Similarly, if the validity test returns a higher likelihood of all sky (case 3) and the value of  $\sigma_{avg}$  is below a given threshold (set at 0.2), the system goes into a “nose-down” mode. Updating of the time-dependent statistical models and  $\sigma_{avg}$  in nose-down mode is the same as in pull-up mode, except that the default update value for  $\sigma_{avg}$  is 0.01 instead of 1.01.

### 3.2 Error detection in horizon estimates

Extreme attitude detection can also help us to detect possible errors in the horizon estimation algorithm; such errors can occur when transient noise causes video degradation. Consider, for example, the following possibility: the validity test returns case 2 (all ground), but  $\sigma_{avg} < 0.8$ . In this situation, we must assume an error occurred in horizon detection, because the aerodynamic characteristics of the plane do not permit such sharp changes in pitch over 1/30th of a second. More generally, if the validity test returns any of the non-normal cases (2, 3 or 4) and the value of  $\sigma_{avg}$  does not conform to the appropriate threshold values, we consider the horizon detection for that frame to be in error. In this case, the horizon estimate from the previous frame is used to estimate the horizon parameters for the current frame.

From extensive flight testing, we observed qualitatively that this extreme attitude and error detection system performs well. It is difficult to quantitatively assess the performance of the system on real-time data since there is no “correct” answer with which to compare it. Both the qualitative viewing of the output, however, along with successful flight tests indicate that the system performs adequately.

Initial tests of the system were performed by playing previously recorded video of human-piloted MAV flights into the system. Figure 2 shows a sample of recorded output data from the system in which the video signal is corrupted by noise. The plot of pitch percentage estimates are overlaid onto the plot of the error detection trigger. When an error is detected, the error detection trigger has a value of 1. As can be seen in Figure 2, the algorithm uses the

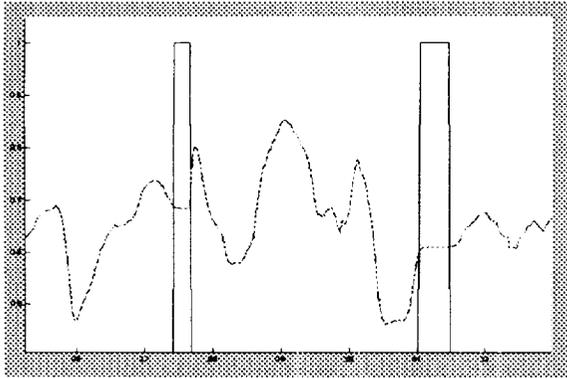


Fig. 2: Pitch percentage estimates with video noise.

last known good estimate during the time of the video transmission error, instead of the estimates from the bad video frames.

Figure 3 illustrates an example of extreme attitude detection. As soon as the horizon disappears off the top of the image, the error detection trigger value jumps to 1. The controller then goes into pull-up mode until the horizon reappears near the top of the image. The pitch estimate remains at 1.01 until the controller exits the pull-up mode.

### 3.3 Kalman filtering

In order to make the horizon estimates usable for self-stabilization and control, the horizon estimates, after being processed by the extreme attitude and error detector, are passed through a Kalman filter [1]. The Kalman filter provides an optimal estimate of a system's current state, given a dynamic system model, a noise model, and a time series of measurements. While a dynamic model of the system is desirable, the formalism of the Kalman filter can be employed even without an accurate dynamic model. Since no dynamic model is readily available for our flexible-wing MAVs<sup>1</sup>, we model the system state (the two parameters of the horizon estimate) as two simple first-order, constant-velocity systems. As such, the Kalman filter has the effect of removing high frequency noise from the system measurements and eliminating any radical single frame errors not first caught by the error detection system. The principal benefit of the Kalman filter for our application is that it effectively eliminates unnecessary small control surface deflections due to noise.

### 3.4 Feedback control

To date, we have employed a very simple controller to validate vision-based flight stability and control for MAVs. For simplicity, the bank angle  $\phi$  and pitch percentage  $\sigma$  are treated as independent from one another, and for both parameters, we implement a simple PD (proportional/derivative) feedback control loop, with gains determined experimentally from flight tests; each control loop is updated at full frame rate (i.e. 30 Hz). In initial flight tests, the derivative gains were set to zero.

1. A dynamic model for our MAV airframes is currently being developed at NASA Langley Research Center. With the code infrastructure in place, the dynamic model can be added when it becomes available.

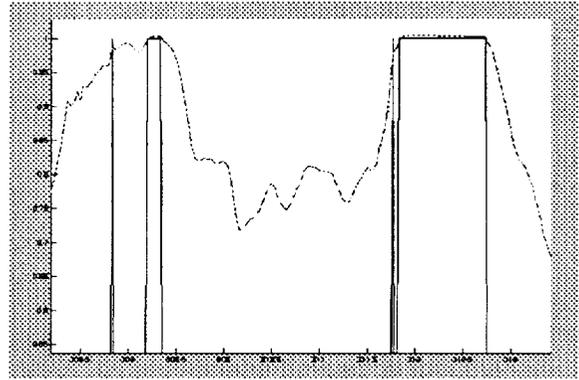


Fig. 3: Extreme attitude detection.

## 4. Self-stabilized flight

### 4.1 Experimental setup

Figure 4 illustrates our current experimental setup. The video signal from the MAV is transmitted from the plane through an antenna to a ground-based computer, where all vision processing is performed. In manual mode, the plane is controlled during flight by a remote human pilot through a standard radio link. In autonomous mode, the plane is controlled through the feedback controller which sends control surface commands to the MAV through a custom designed interface over the same radio link. Our interface allows the PC to control a standard Futaba radio transmitter through an RS-232 serial port.

The MAV used for test flights had a wingspan of 18 inches. While we have designed and flown MAVs with wing spans as small as seven inches, we selected the somewhat larger platform both for its increased dynamic time constants and its ability to carry a high-powered video transmitter (i.e. increased payload). Figure 5 shows a picture of this MAV. The on-board camera is a monolithic CMOS type camera with a 1/3 inch sensor area, and is

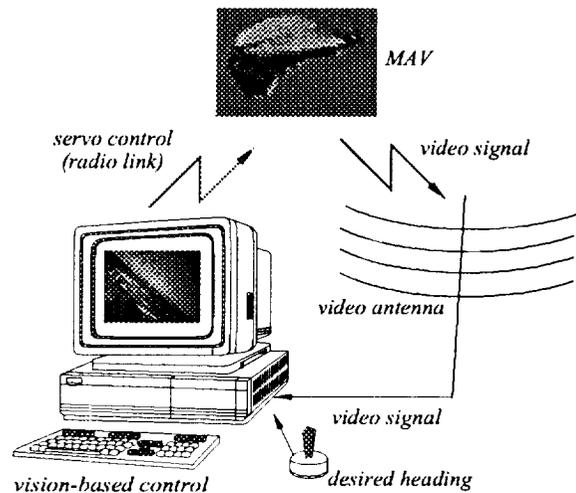


Fig. 4: Experimental setup.

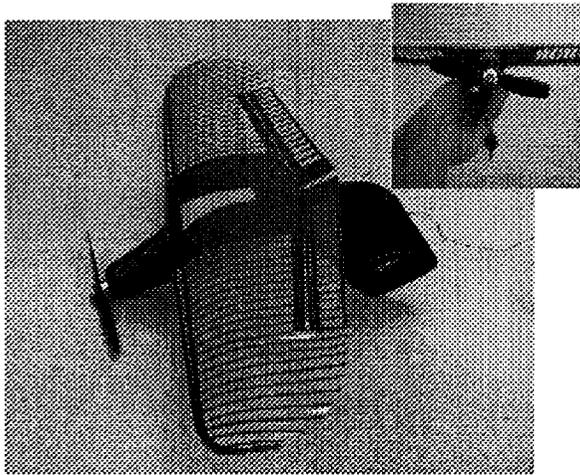


Fig. 5: MAV used for test flights.

connected to an 80 mW video transmitter. The MAV is powered by electric propulsion and has differential elevons for control, although the software is written to support both elevon and rudder-elevator control designs.

The PC interface uses a PIC microcontroller to translate serial commands from the PC into the pulse width modulated signals required for input to the transmitter. A carbon fiber housing was constructed to hold the circuit board and port connectors for the interface.

#### 4.2 Flight testing

Flight testing proceeds as follows. Prior to launch, the aircraft is oriented such that the horizon is in the field-of-view of the camera. This allows the algorithm to build initial models of the sky and the ground; while these models are not used in the horizon-detection algorithm itself [5], they are used for extreme attitude and error detection.

Upon launch, flights are controlled by a human pilot until the MAV reaches sufficient altitude. At that point, control is trans-

ferred to the automated flight control and stability system; in case of catastrophic failure (loss of video signal, etc.), the radio transmitter is equipped with an override button to allow the human pilot to regain control at any time if necessary.

A joystick connected to the PC can be used to adjust the desired heading for the controller. The joystick input effectively commands a bank and pitch angle for the aircraft to follow. Later flights used a pre-programmed set of maneuvers for truly autonomous flight. To date, we have flown uninterrupted autonomous flights of over 10 minutes, flights that ended only due to video transmission interference, or low on-board battery power.

Figure 6 below illustrates a short, 90-second run of actual flight data, where the flight vehicle was under vision-guided control above the University of Florida campus.<sup>1</sup> More recently, the same vision-based control system successfully flew over substantially different terrain at Fort Campbell, Kentucky. Videos of the current algorithm performance under manual and automated flight stabilization can be viewed at <http://mil.ufl.edu/~nechyba/mav>.

#### 5. Conclusion

Qualitatively, even our simple PD control system provides much more stable control than that of our best human pilots, both in terms of steady, level flight, and in coordinated turns. Human pilots can typically not hold the plane on a steady, level heading for more than a few fractions of a second; under vision-guided control, however, we were able to fly long straight segments that were limited only by the range of the video transmitter (see Figure 7 for a 7.5 second stretch of self-stabilized straight and level flight). Prior to the development of the horizon-tracking control system, only pilots with extensive training could learn to fly our micro air vehicles; with the automated control system, however, people who had never piloted any aircraft before were able to easily guide the MAV above the flying arena. It is this fact alone that speaks the most to the potential value of this work. Ideally, one

1. The data in Figure 6(b) appears more oscillatory than the actual flight, because of light non-linear coupling between pitch percentage and roll angle.

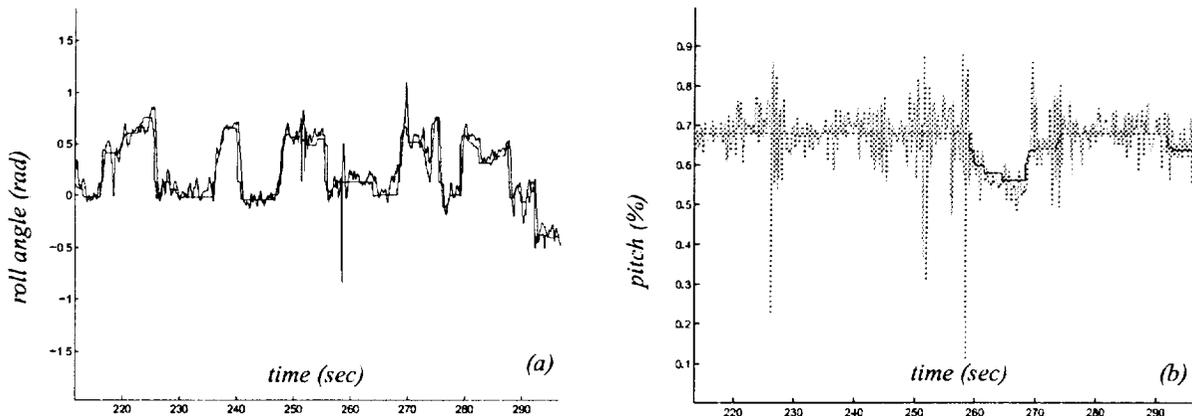
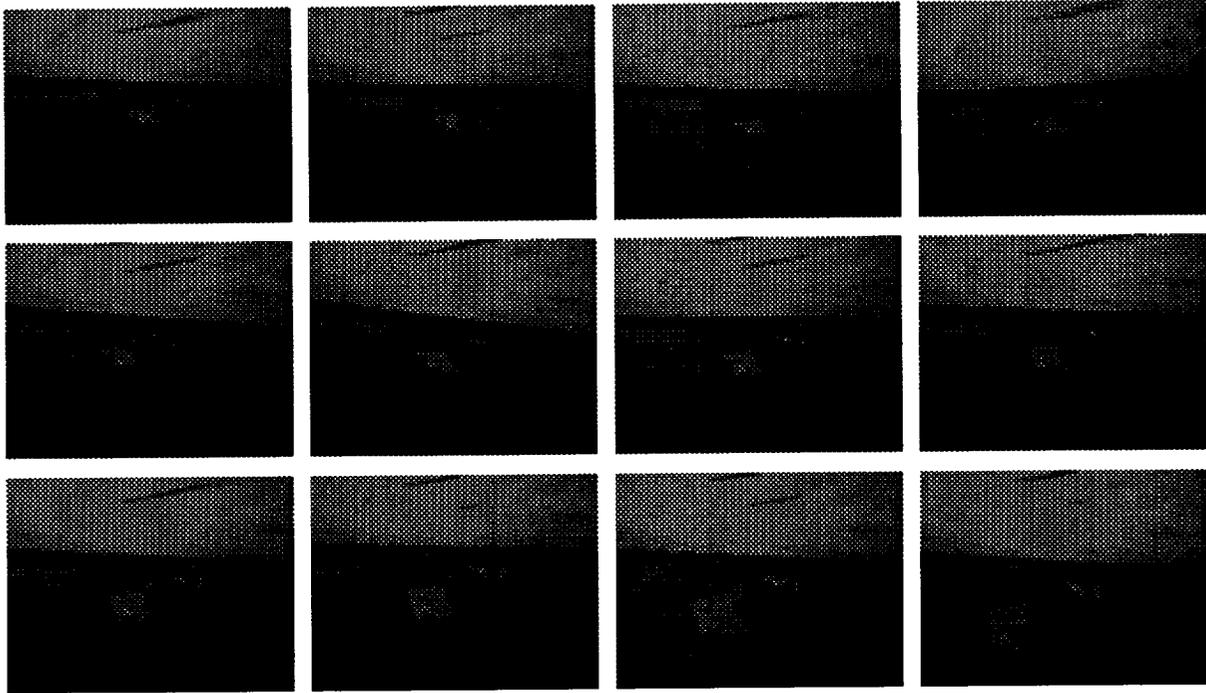


Fig. 6: 90 seconds of sample flight data: (a) commanded roll angle (red) vs. vision-based measured roll angle (blue); (b) commanded pitch percentage (red), vs. vision-based measured pitch percentage.



**Fig. 7: Image sequence of self-stabilized straight and level flight (7.5 seconds, images are 20 frames apart). Note that the black lines at the top of the image are the propeller.**

wants MAVs to be deployable by a wide range of people, not only expert RC pilots; while much remains to be done, including automating landings and take-offs, the work in this paper is a big step towards the development and deployment of usable and practical MAVs.

## 6. References

- [1] B. D. O. Anderson and J. B. Moore, *Optimal Filtering*, Prentice-Hall, Englewood Cliffs, 1979.
- [2] P. R. Ehrlich, D. S. Dobkin and D. Wheye, "Adaptions for Flight," <http://www.stanfordalumni.org/birdsite/text/essays/Adaptions.html>, June, 2001.
- [3] P. R. Ehrlich, D. S. Dobkin and D. Wheye, "Flying in Vee Formation," [http://www.stanfordalumni.org/birdsite/text/essays/Flying\\_in\\_Vee.html](http://www.stanfordalumni.org/birdsite/text/essays/Flying_in_Vee.html), June, 2001.
- [4] S. M. Ettinger, *Design And Implementation Of Autonomous Vision-Guided Micro Air Vehicles*, M.S. Thesis, Electrical and Computer Engineering, University of Florida, August, 2001.
- [5] S. M. Ettinger, M. C. Nechyba, P. G. Ifju and M. Waszak, "Towards Flight Autonomy: Vision-Based Horizon Detection for Micro Air Vehicles," submitted to *Proc. IEEE Int. Conf. on Robotics and Automation*, Washington, D.C., 2002.
- [6] R. Fox, S. W. Lehmkuhle and D. H. Westendorf, "Falcon Visual Acuity," *Science*, vol. 192, pp. 263-5, 1976.
- [7] P. G. Ifju, S. Ettinger, D. A. Jenkins and L. Martinez, "Composite Materials for Micro Air Vehicles," *Proc. of the SAMPE Annual Conf.*, Long Beach CA, May 6-10, 2001.
- [8] P. G. Ifju, S. Ettinger, D. A. Jenkins, and L. Martinez, "Composite Materials for Micro Air Vehicles," to appear in *SAMPE Journal*, July 2001.
- [9] D. A. Jenkins, P. Ifju, M. Abdulrahim and S. Olipra, "Assessment of Controllability of Micro Air Vehicles," *Proc. Sixteenth Int. Conf. On Unmanned Air Vehicle Systems*, Bristol, United Kingdom, April 2001.
- [10] J. Kellogg, et. al., "The NRL Mite Air Vehicle," *Proc. Sixteenth Int. Conf. On Unmanned Air Vehicle Systems*, pp. 25.1-25.14, Bristol, United Kingdom, April 2001.
- [11] Northern Prairie Wildlife Research Center, "Migration of Birds: Orientation and Navigation," <http://www.npwrc.usgs.gov/resource/othrdata/migration/ori.htm>, June, 2001.
- [12] M. Ollis and T. Stentz, "Vision-Based Perception for an Automated Harvester," *Proc. IEEE Int. Conf. on Intelligent Robots and Systems*, vol. 3, pp. 1838-44, 1997.
- [13] G. Ritchison, "Ornithology: Nervous System: Brain and Special Senses II," <http://www.biology.eku.edu/RITCHISO/birdbrain2.html>, June 2001.
- [14] W. Shyy, D. A. Jenkins and R. W. Smith, "Study of Adaptive Shape Airfoils at Low Reynolds Number in Oscillatory Flows," *AIAA Journal*, vol. 35, pp.1545-48, 1997.
- [15] R. W. Smith and W. Shyy, "Computation of Aerodynamics Coefficients for a Flexible Membrane Airfoil in Turbulent Flow: A Comparison with Classical Theory," *Phys. Fluids*, vol. 8, no. 12, 1996.
- [16] G. C. Whittow, ed., *Sturkie's Avian Physiology, Fifth Ed.*, Academic Press, San Diego, 2000.